

Online Model-Free Influence Maximization with Persistence

Paul Lagr  e

LRI, Univ. Paris-Sud, Univ. Paris-Saclay
paul.lagree@u-psud.fr

Bogdan Cautis

Noah’s Ark Lab, Huawei, Hong Kong
cautis.bogdan@huawei.com

Olivier Capp  

LIMSI, CNRS, Univ. Paris-Saclay
cappe@limsi.fr

Silviu Maniu

LRI, Univ. Paris-Sud, Univ. Paris-Saclay
silviu.maniu@u-psud.fr

ABSTRACT

Influence maximization is the problem of finding influent nodes in a graph so as to maximize the spread of information. It has many applications in advertising and marketing on social networks. In this paper, we study the problem of sequentially selecting seeds in the network under the hypothesis that previously activated nodes can still transfer information, but do not yield further rewards. Furthermore, we make no assumption on the underlying diffusion model. We refer to this problem as *online influence maximization with persistence*. We first discuss scenarios motivating the problem and present our approach to solve it. We then analyze a novel algorithm relying on upper confidence bound on the so-called *missing mass*, that is, the expected number of nodes that can still be reached from a given seed. From a computational standpoint, the proposed approach is several orders faster than state-of-the-art methods, making it possible to tackle very large graphs. In addition, it displays high-quality spreads on both simulated and real datasets.

CCS CONCEPTS

•Information systems → Data mining; •Computing methodologies → Online learning settings;

KEYWORDS

Influence Maximization, Online Learning, Multi-Armed Bandits

1 INTRODUCTION

Advertising based on word-of-mouth diffusion in social media has become very important in the digital marketing landscape. Nowadays, social value and social influence are arguably the hottest concepts in the area of Web advertising and most companies that advertise in the Web space must have a “social” strategy. For example, on widely used platforms such as Facebook or Twitter, promoted posts are interleaved with normal posts on user feeds. Users interact with these posts by actions such as “likes” (adoption), “shares” or “reposts” (network diffusion). This represents an unprecedented tool in advertising, be it with a commercial intent or not, as products, news, ideas, movies, political manifests, tweets, etc. can propagate easily to a large audience [32, 33].

Motivated by the need for effective viral marketing strategies, *influence estimation* and *influence maximization* (IM) have become important research problems, at the intersection of data mining and social sciences [10]. In short, IM is the problem of selecting a set of nodes from a given diffusion graph, maximizing the expected

spread under an underlying diffusion model. This problem was introduced in 2003 by the seminal work of Kempe et al. [18], through two stochastic, discrete-time diffusion models, *Linear Threshold* (LT) and *Independent Cascade* (IC). These models rely on diffusion graphs whose edges are weighted by a score of influence. They show that selecting the set of nodes maximizing the expected spread is NP-hard for both models, and they propose a greedy algorithm that takes advantage of the sub-modularity property of the influence spread, but does not scale to large graphs. A rich literature followed, focusing on computationally efficient and scalable algorithms to solve IM. The recent benchmarking study of Arora et al. [1] summarizes state-of-the-art techniques and also debunks many IM myths. In particular, it shows that, depending on the underlying diffusion model and the choice of parameters, each algorithm’s behavior can vary significantly, from very efficient to prohibitively slow.

Importantly, all the IM studies discussed in [1] have as starting point a specific model (IC or LT), whose graph topology and parameters – basically the edge weights – are known. However, this assumption is unrealistic, and, recently, we have witnessed a trend towards bridging the gap between theory and practical relevance in the IM framework, along several dimensions.

In particular, one such dimension is the one of *offline, model-specific methods*, which can *infer* the diffusion parameters or the underlying graph structure (if unknown, or, as often the case, implicitly overlaying the existing social graph), or both, starting from observed information cascades [9, 11–13, 15, 27]. In short, information cascades are time-ordered sequences of records indicating when a specific user adopted a specific item.

There are however many situations where it is unreasonable or counter-productive to assume the existence of relevant historical data in the form of cascades. For such settings, *online approaches*, which can learn the underlying diffusion parameters *while running diffusion campaigns*, have been proposed. Bridging IM and inference, this is done by balancing between exploration steps (of yet uncertain model aspects) and exploitation ones (of the best solution so far), by so called *multi-armed bandits* techniques, where an agent interacts with the network to infer influence probabilities [8, 30, 34]. The learning agent sequentially selects seeds from which diffusion processes are initiated in the network; the obtained feedback is used to update the agent’s knowledge of the model.

Nevertheless, all these studies on inferring diffusion networks, whether offline or online, rely on parametric diffusion models, i.e., assume that the actual diffusion dynamics are well captured by such a model (e.g., IC). This maintains significant limitations for practical purposes. First, the more complex the model, the harder to learn, especially in campaigns that have a relatively short timespan,

making model inference and parameter estimation very challenging within a small horizon (typically tens or hundreds of spreads). Second, it is commonly agreed that the aforementioned diffusion models represent elegant yet coarse interpretations of a reality that is much more complex and often hard to observe fully. For examples of insights into this complex reality, the *topical* or *non-topical* nature of an influence campaign, the *popularity* of the piece of information being diffused, or its specific *topic* were all shown to have a significant impact on hashtag diffusions in Twitter [9, 16, 25].

Our contribution. Aiming to address such limitations, we study in this paper a *model-free approach for online and adaptive IM*, in which the underlying assumptions for the diffusion processes are kept to a minimal (if, in fact, hardly any). We argue that it can represent a versatile tool in many practical scenarios. More precisely, we focus on social media diffusion scenarios in which influence campaigns consist of multiple *consecutive trials* (or *rounds*) spreading the same piece of information (be it a product, idea, post, hashtag, etc). The goal of each campaign is to reach (or *activate*) as many distinct users as possible, the objective function being the total spread. At each round, the learning agent selects the seeds from which a new diffusion process starts in the network, assuming a certain notion of *persistence*, which is given the following interpretation: (i) during a campaign, the agent may “re-seed” certain nodes (we may want to ask a particular node to initiate spreads several times, e.g., if it has a strong *converting impact*), and (ii) nodes who were already activated in the ongoing campaign, i.e., have adopted that piece of information, remain “committed to the cause” throughout, and thus continue spreading the information and exerting influence on their peers. Political campaigns are one obvious example of such a scenario, but not the only one. Nowadays, many online marketing strategies rely heavily on their adopters being “loyalists” or even “brand fanatics” (e.g., think of Disney’s latest Star Wars movies).

We call this problem *online influence maximization with persistence* (OIMP). Our solution for it follows the multi-armed bandit idea initially employed in Lei et al. [19], but we adopt instead a *model-free perspective*, whose only input is population targetted. As an abstraction, our approach represents the graph as a forest of depth-1 trees which are key to build the statistics that control the decisions of the algorithm. As in [19], we assume that different campaigns are independent, but we further simplify the model by the fact that the only feedback the agent can gather after each trial are the activated nodes. The rationale is that oftentimes, for a given “viral” item, we can track in applications only *when* it was adopted by various users, but not *why*. In particular, our approach does not require the observation of successful or failed edge activations. In our setting, a key difference w.r.t. other multi-armed bandit studies for IM such as [8, 30, 34] is that these look for a *constant* optimal set of seeds, while the difficulty with OIMP is that the seemingly best action at a given trial depends on the activations of the previous trials (and thus the learning agent’s past decisions).

The multi-armed bandit algorithm we propose, called GT-UCB, relies on a famous statistical tool known as the *Good-Turing estimator*, first developed during WWII to crack the Enigma machine, and later published by Good in a study on species discovery [14]. Our approach is inspired by the work of Bubeck et al. [6], which proposed the use of the Good-Turing estimator in a context where

the learning agent needs to sequentially select experts that only sample one of their child nodes at each trial. In contrast, in OIMP, when a seed node is selected, it may have a potentially large spread and may activate many nodes at once. Our solution follows the well-known *optimism in the face of uncertainty* principle from the bandit literature (see [5] for an introduction to multi-armed bandit problems). In short, we derive an upper-confidence bound on the estimator for the remaining spread potential of each seed, and we choose in a principled manner between explore and exploit steps.

Through this approach, we show that efficient and effective influence maximization and diffusion campaigns can be done in a highly uncertain or under-specified social environment, along with formal guarantees on the spread.

The paper is organized as follows. We review other related works in Section 2 and we formalize the IM model in Section 3. In Section 4, we present the GT-UCB algorithm and we provide a theoretical analysis for it in Section 5. We evaluate empirically this algorithm in Section 6, by comparison to the state-of-the-art method from [19], on both simulated and real data. We conclude in Section 7.

2 OTHER RELATED WORK

We have already discussed in Section 1 some of the main related studies in the area of IM. For further details, we refer the interested reader to the recent survey in [1], which discusses the pros and cons of the best known techniques for IM. In particular, the authors highlight that the *Weighted Cascade* (WC) instance of IC, where the weights associated to a node’s incoming edges must sum to one, leads to poor performance for otherwise rather fast IC algorithms. They conclude that PMC [24] is the state-of-the-art method to efficiently solve the IC optimization problem, while TIM+ [29] and IMM [28] – later improved by [23] with SSA – are the best current algorithms for WC and LT models.

Besides the already discussed offline methods for inferring the diffusion network and its parameters, we mention here that a first *offline and model-free method* for inferring the diffusion network from existing cascades has been proposed recently in [26]. We have in common with this work the goal to devise generic, non-parametric methods, yet in a online IM framework.

Other methods have been devised to handle the prevalent uncertainty in diffusion media, e.g., when replacing edge probability scores with ranges thereof, by solving an IM problem whose *robust* outcome should provide some effectiveness guarantees w.r.t. all possible instantiations of the uncertain model [7, 17].

Methods for IM that take into account more detailed information, such as topical categories, have been considered in the literature [2, 9, 31]. Interestingly, [25] experimentally validates the intuition that different kinds of information spread differently in social networks, by relying on two complementary properties, namely *stickiness* and *persistence*. The former can be seen as a measure of how viral the piece of information is, passing from one individual to the next. The latter can be seen as an indicator of the extent to which repeated exposures to that piece of information impact its adoption, and it was shown to characterize *complex contagions*, of controversial information (e.g., from politics). Their notion of persistence is related to a certain degree to ours, since it pertains to repeated exposures to items in one’s social circle and their effects on spread.

3 MODEL

The goal of the *online influence maximization with persistence* is to successively select (or *activate*) a number of seed nodes in the diffusion graph, in order to *reach* (or *spread* to) as many other nodes as possible. We formally define this problem next.

3.1 Background and problem definition

The traditional problem of influence maximization (IM) is to select a set of seed nodes I , under a cardinality constraint $|I| = L$, such that the expected *spread* of an influence cascade starting from I (or the expected number of activated nodes) is maximized. Formally, denoting by the random variable $S(I)$ the spread initiated by the seed set I , IM aims to solve the following optimization problem:

$$\arg \max_{I \subseteq V, |I|=L} \sigma(I) := \mathbb{E}[S(I)].$$

As mentioned before, a plethora of algorithms have been proposed to solve the IM problem, under specific diffusion models. These algorithms can be viewed as *full-information* and *offline* approaches: they choose all the seeds at once, in one step, and they have the complete diffusion configuration, i.e., the graph topology and the influence probabilities.

In the *online* case, during a sequence of N (what we call hereafter the *budget*) consecutive trials, L seed nodes are selected at each trial, and *feedback* on the achieved spread from these seeds is collected. Without precise influence mechanisms, the diffusion process in the graph G can be loosely described as follows:

Definition 3.1 (Influence cascade). Given a graph $G = (V, E)$, by selecting and activating a set of seed nodes I_t at step t in a campaign, a diffusion is initiated and eventually terminates, reaching and activating in the process other nodes in G . The observed *feedback* (or *spread*) $S(I_t)$ for I_t 's selection are all the activated nodes, while the associated *reward* consists of only the *newly activated* ones.

The selection at a given trial can take into account (adapt to) the feedback obtained at the previous trials. Formally, the problem becomes the following:

PROBLEM 1 (OIMP). Given a graph $G = (V, E)$, a budget of N trials, and a number $1 \leq L \leq |V|$ of nodes to be activated at each trial, the objective of the online influence maximization with persistence (OIMP) is to solve the following optimization problem:

$$\arg \max_{I_n \subseteq V, |I_n|=L, \forall 1 \leq n \leq N} \mathbb{E} \left[\bigcup_{1 \leq n \leq N} S(I_n) \right].$$

As noticed in [19], the traditional offline IM can be seen as a special instance of the online one, where the budget is $N = 1$. Note that, in contrast to persistence-free online influence maximization – considered, e.g., in [30, 34] – the performance criterion used in OIMP displays the so-called *diminishing returns property*: the expected number of nodes activated by successive selections of a given seed is decreasing, due to the fact that nodes that have already been activated are discounted. We refer to the expected number of nodes remaining to activate as the *potential* or *missing mass* of a seed. The diminishing returns property implies that there is no static best set of seeds to be selected, but that the algorithm must follow an adaptive policy, which can detect that the remaining




Figure 1: Forest of 3 depth-1 experts.

potential of a seed is small and switch to another seed that has been less exploited. Our solution to this problem has to overcome challenges on two fronts: (1) it needs to estimate the potential of nodes at each round, without knowing the diffusion model nor the activation probabilities, and (2) it needs to identify the currently best seeds, according to their estimated potentials.

Other approaches for the online IM problem rely on estimating diffusion parameters [19, 30, 34] – generally, a distribution over the influence probability of each edge in the graph. However, the assumption that one can estimate accurately the diffusion parameters – and notably the diffusion probabilities – may be overly ambitious, especially in cases where the number of allowed trials (the budget) is rather limited. A limited trial setting is arguably more in line with real-world campaigns: take as example political or marketing campaigns, which only last for a few days.

In our approach, we work with parameters on *nodes*, instead of edges. More specifically, these parameters represent the potentials of remaining spread from each of the candidate seed nodes. We stress that this potential can evolve as the campaign proceeds. In this way, we can go around the dependencies on specific diffusion models, and furthermore, we can relax significantly the dependency on a detailed graph topology, as discussed next.

3.2 From general graphs to forests of experts

By design, the internal graph topology on which our OIMP solution executes is a *forest of depth-1 trees*, where the root of each tree represents an *expert* – in our setting, the experts are the candidates for seed selection. Each expert is connected to an unknown and potentially large base (the expert's *support*) of basic nodes, each with an unknown activation probability. For illustration, we give in Figure 1 an example of such a depth-1 forest, having 3 experts connected to 4, 3, and 4 basic nodes, respectively.

We stress that the input for our method remains a general diffusion graph, as defined in Problem 1, and this represents the medium over which at each trial the real diffusions happen and spread can be observed. However, the simpler, depth-1 forest representation is the one on which the sequence of adaptive IM trials of the campaign will perform seed selections.

We thus complete the formal setting by assuming the existence of K expert nodes in the diffusion graph $G = (V, E)$ – corresponding to the number of trees in the depth-1 forest – such that each expert $k \in [K]$ is connected to a set $A_k \subseteq V$ of basic nodes. We denote $p_k(u)$ the probability for expert k to activate the child node $u \in A_k$.

In the context of forests of depth-1 trees, the diffusion process can be abstracted as follows:

Definition 3.2 (Depth-1 forest influence cascade). When an expert $k \in [K]$ is selected at a given trial, each basic node $u \in A_k$ is sampled for activation, according to its probability $p_k(u)$. The *feedback* (or *spread*) for k 's selection consists of all the activated nodes, while the associated *reward* consists of the *newly activated* ones.

Remark. Limiting the adaptive online method to working internally with such a forest of depth-1 trees may seem overly restrictive, but our thesis is that this abstraction generalizes very well to the current real-life social network scenarios. First, despite the fact that we model the reach of every influencer by 1-hop links to the to-be-influenced nodes, these edges may represent longer paths in the real graph G . Second, since large social graphs are known to have a modular structure, it is reasonable to assume that, by choosing the experts carefully, one can maintain a minimal overlap in their scopes of influence. Indeed, in social networks composed of densely-connected communities with relatively few connections to other communities, big hubs do not usually link to each other, and those that are far from each other do not influence the same region of the graph. Moreover, since most social networks are scale-free, with a few large hubs controlling the information flow through the entire network, they exhibit good candidates for the role of experts. In Section 4.2, we describe several strategies for choosing experts that are as independent as possible, in this way abstracting away from a general graph topology which may not fully describe the diffusion channels anyway.

3.3 Missing mass and Good-Turing estimator

Given the K experts, the OIMP problem boils down to the following: *How should we select an expert at each step?* More precisely, a good algorithm for OIMP should aim at selecting the expert k with the largest potential for influencing its children A_k . However, the true potential value of an expert is *a priori* unknown to the decision maker. We now describe our approach to estimate this value, using the concept of *missing mass*:

Definition 3.3 (Missing mass $R_{k,n}$). Consider an expert $k \in [K]$ connected to A_k basic nodes and having already initiated influence cascades in n trials (not necessarily consecutive), with respective feedback $S_{k,1}, \dots, S_{k,n}$, where each $S_{k,i} \subseteq V$ is the set of nodes that were activated. The *missing mass* $R_{k,n}$ is the expected number of *new* nodes that would be activated upon starting a $n+1$ th cascade from k :

$$R_{k,n} := \sum_{u \in A_k} \mathbb{1} \left\{ u \notin \bigcup_{i=1}^n S_{k,i} \right\} p_k(u),$$

where $\mathbb{1}\{\cdot\}$ denotes the indicator function.

Definition 3.3 provides a formal way to obtain the remaining potential of an expert k at a given time. The difficulty is, however, that the probabilities $p_k(u)$ are unknown. Hence, we have to design a *missing mass estimator* $\hat{R}_{k,n}$ instead. It is important to stress that the missing mass is a random quantity, because of the dependency on the spreads $S_{k,1}, \dots, S_{k,n}$. Due to the diminishing returns property, the sequence $(S_{k,n})_{n \geq 1}$ is stochastically decreasing.

Following ideas from [6, 14], we now introduce a version of the Good-Turing statistic, tailored to our problem of estimating the missing mass. For the sake of clarity, we omit in the following the subscript denoting the expert k . Let S_1, \dots, S_n be the n cascades

sampled independently from an expert. We denote by $U_n(u)$ the binary function whose value is 1 if node u has been activated *exactly* once – such occurrences are called *hapaxes* in linguistics. The idea of the Good-Turing estimator is to estimate the missing mass as the proportion of hapaxes in the n sampled cascades, as follows:

$$\hat{R}_n := \frac{\sum_{u \in A} U_n(u)}{n}.$$

Albeit simple, this estimator turns out to be quite effective in practice. If an expert is connected to a combination of nodes having high activation probabilities, along with nodes having low activation probabilities, then successive cascades sampled from this expert will result in multiple activations of the high-probability nodes and few of the low-probability ones. Hence, after a few cascades, the expert's potential will be low, a fact that will be captured by the low proportion of hapaxes.

Remark. While bearing similarities with the traditional missing mass concept, we highlight one fundamental difference of our problem w.r.t the one studied in [6], which impacts both the algorithmic solution and the analysis. Since at each step, after selecting an expert, *every* node connected to that expert is sampled, the algorithm receives a more general (larger) feedback than in [6], whose feedback is in $[0, 1]$. Interestingly, the quantity $\lambda := \sum_{u \in A} p(u)$, which corresponds to the expected number of basic nodes an expert activates or re-activates in a cascade, will prove to be a crucial ingredient for our problem.

The classic Good-Turing estimator is known to be slightly biased (see Theorem 1 in [21] for example). We show in Lemma 3.4 that our missing mass estimator adds an additional factor λ to this bias:

LEMMA 3.4. *The bias of the missing mass estimator is*

$$\mathbb{E}[R_n] - \mathbb{E}[\hat{R}_n] \in \left[-\frac{\lambda}{n}, 0 \right].$$

PROOF.

$$\begin{aligned} \mathbb{E}[R_n] - \mathbb{E}[\hat{R}_n] &= \sum_{u \in A} \left[p(u)(1-p(u))^n - \frac{n}{n} p(u)(1-p(u))^{n-1} \right] \\ &= -\frac{1}{n} \sum_{u \in A} p(u) \times n p(u)(1-p(u))^{n-1} \\ &= -\frac{1}{n} \mathbb{E} \left[\sum_{u \in A} p(u) U_n(u) \right] \in \left[-\frac{\sum_{u \in A} p(u)}{n}, 0 \right] \quad \square \end{aligned}$$

3.4 Confidence interval for the missing mass

To develop our UCB-like algorithm in Section 4, we need to control the value of our estimator for the missing mass. As shown in Lemma 3.4, in expectation, the estimation should be relatively accurate. However, in order to understand what may happen in the worst-case, we need to characterize the deviation of the Good-Turing estimator:

THEOREM 3.5. *With probability at least $1 - \delta$, for $\lambda = \sum_{u \in A} p(u)$ and $\beta_n := \left(1 + \sqrt{2}\right) \sqrt{\frac{\lambda \log(4/\delta)}{n}} + \frac{1}{3n} \log \frac{4}{\delta}$, the following holds:*

$$-\beta_n - \frac{\lambda}{n} \leq R_n - \hat{R}_n \leq \beta_n.$$

Note that the additional term appearing in the left deviation corresponds to the bias of our estimator, which leads to a non-symmetrical interval.

PROOF. We prove the confidence interval in three steps: (1) Good-Turing estimator deviation, (2) missing mass deviation, (3) combining the previous two inequalities for the final confidence interval.

Note that the samples of different nodes are assumed *independent*. This is a simplification with respect to the classic missing mass concentration results, which rely on negative association [20, 21]. On the other hand, since we may activate several nodes at once, we need concentration bounds to control the increments of R_n .

(1) **Good-Turing deviations.** Recall that $\hat{R}_n = \sum_{u \in A} \frac{U_n(u)}{n}$. For $X_n(u) := \frac{U_n(u)}{n}$, we have next that

$$\begin{aligned} v &:= \sum_{u \in A} \mathbb{E}[X_n(u)^2] = \frac{1}{n^2} \sum_{u \in A} \mathbb{E}[U_n(u)] \\ &= \frac{1}{n^2} \sum_{u \in A} np(u)(1-p(u))^{n-1} \leq \frac{\lambda}{n}. \end{aligned}$$

Moreover, clearly the following holds: $X_n(u) \leq \frac{1}{n}$.

Applying Bennett's inequality (Theorems 2.9, 2.10 in [4]) to the independent random variables $\{X_n(u)\}_{u \in A}$ yields

$$\mathbb{P}\left(\hat{R}_n - \mathbb{E}[\hat{R}_n] \geq \sqrt{\frac{2\lambda \log(1/\delta)}{n}} + \frac{1}{3n} \log(1/\delta)\right) \leq \delta. \quad (1)$$

The same inequality can be derived for left deviations.

(2) **Missing mass deviations.** Let $Z_n(u)$ denote the indicator equal to 1 if u has never been activated up to trial n . We can rewrite the missing mass as follows:

$$R_n = \sum_{u \in A} Z_n(u)p(u).$$

Let $Y_n(u) = p(u)(Z_n(u) - \mathbb{E}[Z_n(u)])$ and $q(u) = \mathbb{P}(Z_n(u) = 1) = (1-p(u))^n$. For some $t > 0$, we have next that

$$\begin{aligned} \mathbb{P}(R_n - \mathbb{E}[R_n] \geq \epsilon) &\leq e^{-t\epsilon} \prod_{u \in A} \mathbb{E}\left[e^{tY_n(u)}\right] \\ &= e^{t\epsilon} \prod_{u \in A} \left(q(u)e^{tp(u)(1-q(u))} + (1-q(u))e^{-tp(u)q(u)}\right) \\ &\leq e^{-t\epsilon} \prod_{u \in A} \exp(p(u)t^2/(4n)) = \exp\left(-t\epsilon + t^2/(4n)\lambda\right). \end{aligned}$$

The first inequality is well-known in exponential concentration bounds and relies on Markov's inequality. The second inequality follows from [3] (Lemma 7) (for the reader's convenience, also recalled as Lemma A.2 in the Appendix).

Then, choosing $t = \frac{2n\epsilon}{\lambda}$, we obtain

$$\mathbb{P}\left(R_n - \mathbb{E}[R_n] \geq \sqrt{\frac{\lambda \log(1/\delta)}{n}}\right) \leq \delta. \quad (2)$$

We can proceed similarly to obtain the left deviation.

(3) **Putting it all together.** We can combine Lemma 3.4 with Eq. (1), (2), to obtain the final result. Note that we need to replace δ by $\frac{\delta}{4}$ to ensure that both the left and right bounds for the Good-Turing estimator and the missing mass are verified. \square

4 ALGORITHM

In this section, we describe our UCB-like algorithm, which relies on the Good-Turing estimator to sequentially select the experts (seeds) to activate at each round.

4.1 Upper confidence bounds

Following principles from the bandit literature, the GT-UCB algorithm relies on *optimism in the face of uncertainty*. At each step (trial) t , the algorithm selects the highest upper-confidence bound on the missing mass – denoted by $b_k(t)$ in the following – and activates (plays) the corresponding expert k . This algorithm achieves robustness against the stochastic nature of the cascades, by ensuring that experts who “underperformed” with respect to their potential in previous trials may still be selected later on. Consequently, GT-UCB aims to maintain a degree of *exploration* of experts, in addition to the *exploitation* of the best experts as per the feedback gather so far. By relying on the missing mass estimator, our solution also takes into account one aspect of the *persistence* property: the diminishing returns in efficiency of experts that have already been activated many times.

Algorithm 1 – GT-UCB ($L = 1$)

Require: Set of experts $[K]$, time budget N , (unknown) diffusion graph G

- 1: **Initialization:** play each expert $k \in [K]$ once, observe the spread $S_{k,1}$ in G , set $n_k = 1$
 - 2: For each $k \in [K]$: update the reward $W = W \cup S_{k,1}$
 - 3: **for** $t = K + 1, \dots, N$ **do**
 - 4: Compute $b_k(t)$ for every expert k
 - 5: Choose $k(t) = \arg \max_{k \in [K]} b_k(t)$
 - 6: Play expert $k(t)$ and observe spread $S(t)$ in G
 - 7: Update cumulative reward: $W = W \cup S(t)$
 - 8: Update statistics of expert $k(t)$: $n_{k(t)}(t+1) = n_{k(t)}(t) + 1$ and $S_{k(t),n_{k(t)}(t)} = S(t)$.
 - 9: **end for**
 - 10: **return** W
-

Algorithm 1 presents the main components of GT-UCB for the case $L = 1$, that is, when a single expert is chosen at each step.

The algorithm starts by activating each expert $k \in [K]$ once, in order to initialize its Good-Turing estimator. The main loop of GT-UCB occurs at lines 3-9. Let $S(t)$ be the observed spread at the trial t , and let $S_{k,s}$ be the result of the s -th diffusion initiated at expert k . At every step $t > K$, we recompute for each expert $k \in [K]$ its index $b_k(t)$, representing the upper confidence bound on the expected reward in the next trial. The computation of this index uses the previous samples $S_{k,1}, \dots, S_{k,n_k(t)}$ and the number of times each expert k has been activated up to trial t , $n_k(t)$. Using the result in Theorem 3.5, with confidence probability $\delta = \frac{1}{t}$, the upper confidence bound can be computed as follows:

$$b_k(t) = \hat{R}_k(t) + \left(1 + \sqrt{2}\right) \sqrt{\frac{\hat{\lambda}_k(t) \log(4t)}{n_k(t)}} + \frac{1}{3n_k(t)} \log(4t),$$

where $\hat{R}_k(t)$ is the Good-Turing estimator and $\hat{\lambda}_k(t) := \sum_{s=1}^{n_k(t)} \frac{|S_{k,s}|}{n_k(t)}$ is an estimator for the expected spread from expert k .

Then, in line 5, GT-UCB selects the expert $k(t)$ with the largest index, and initiates a cascade from this node. The feedback $S(t)$ is observed in the diffusion graph G and is used to update the cumulative reward set W . $S(t)$ returns only the ids of the nodes being activated, with no information on *how* this diffusion happened in G ; indeed, the diffusion may use the persistence property, i.e.,

that the newly activated nodes come from activation paths in G , passing through nodes activated at a trial $t' < t$. Finally, statistics associated to the chosen expert $k(t)$ are updated.

4.2 Extracting experts

GT-UCB does not make any assumptions about the topology of the graph G from which the experts come. Indeed, in some settings it may be more natural to assume that the set of experts is given and that the activations at each trial can be observed, while G 's connections remain unknown. In other settings, we may start from an existing social network G , in which case we need to extract a set of K experts from it. Recall that, for our method to be effective, ideally, we should choose experts that have as little intersection as possible in their "scopes of influence". While this may be interpreted and performed differently, from one application to another, we discuss next some of the most natural heuristics for selecting experts.

MaxDegree. This method selects the K nodes with the highest out-degrees in G . Note that by this criterion we may select experts with overlapping influence scopes.

Greedy MaxCover. This strategy follows the well-known greedy approximation algorithm for selecting a cover of the graph G . Specifically, the algorithm executes the following steps K times:

- (1) Select the node with highest out-degree
- (2) Remove all out-neighbors of the selected node

To limit intersections among expert scopes even more, nodes reachable by more than 1 hops may be removed at step (2).

DivRank [22]. DivRank is a PageRank-like method relying on reinforced random walks, with the goal of producing diverse high-ranking nodes, while maintaining the rich-gets-richer paradigm. We adapted the original DivRank procedure by inverting the edge directions. In doing so, we get influential nodes instead of prestigious ones. By selecting the K highest scoring nodes as experts, the diversity is naturally induced by the reinforcement of random walks. This ensures that the experts are fairly scattered in the graph and should have limited impact on each other.

IM approximated algorithms. The fourth method we tested in our experiments assigns uniformly at random a propagation probability to each edge of G , assuming the IC model. Then, a state-of-the-art IM algorithm – PMC in our experiments – is executed on G to get the set of K experts having the highest potential spread.

4.3 Extensions for the case $L > 1$

Algorithm 1 can be easily adapted to select $L > 1$ experts at each round. Instead of choosing the expert maximizing the Good-Turing UCB in line 5, we can select the experts having the L largest indices. Note that $k(t)$ becomes now a set of L experts. A diffusion is initiated from the associated nodes in G and, at termination, all activations are observed. In Section 3, we assumed that all sets A_k of basic nodes are disjoint. In practice, even if the experts are fairly scattered in the graph, they may now jointly play a role in some activations. Therefore, we propose a simple heuristic to assign activated nodes to selected experts, by a breadth-first approach, as follows: for an activated node $u \in S(t)$, we assign this node to the selected expert reachable from u by the shortest *live path* in G , where a live path corresponds to a sequence of activated nodes from $S(t)$.

5 THEORETICAL GUARANTEES

In this section, we provide an analysis of the *waiting time* (defined below) of GT-UCB, by comparing it to the waiting time of an oracle policy, following ideas from [6]. Let $R_k(t)$ be the missing mass of expert k at trial number t . Note that this differs from $R_{k,n}$, which is the missing mass of expert k once it has been played n times.

Definition 5.1 (Waiting time). Let $\lambda_k = \sum_{u \in A_k} p(u)$ denote the expected number of activations obtained by the first call to expert k . For $\alpha \in (0, 1)$, the *waiting time* $T_{UCB}(\alpha)$ of GT-UCB represents the round at which the missing mass of *each* expert k is smaller than $\alpha\lambda_k$. Formally,

$$T_{UCB}(\alpha) := \min\{t : \forall k \in [K], R_k(t) \leq \alpha\lambda_k\}.$$

Note that this definition can be applied to any strategy for expert selection and, in particular, to an oracle one that knows beforehand the α value that is targeted, the sampled spreads $(S_{k,s})_{k \in [K], s \geq 1}$, and the individual activation probabilities $p_k(u)$, $u \in A_k$. A policy having access to all these aspects will perform the fewest possible activations on each expert. We denote by $T^*(\alpha)$ the waiting time of the oracle policy. We are now ready to state the main theoretical property of the GT-UCB algorithm.

THEOREM 5.2 (WAITING TIME). Let $\lambda^{\min} := \min_{k \in [K]} \lambda_k$ and let $\lambda^{\max} := \max_{k \in [K]} \lambda_k$. Assuming that $\lambda^{\min} \geq 13$, for any $\alpha \in \left[\frac{13}{\lambda^{\min}}, 1\right]$, if we define $\tau^* := T^*\left(\alpha - \frac{13}{\lambda^{\min}}\right)$, with probability at least $1 - \frac{2K}{\lambda^{\max}}$ the following holds:

$$T_{UCB}(\alpha) \leq \tau^* + K\lambda^{\max} \log(4\tau^* + 11K\lambda^{\max}) + 2K.$$

The proof of this result is given in Appendix B. Unsurprisingly, Theorem 5.2 says that GT-UCB must perform slightly more activations of the experts than the oracle policy. With high probability – assuming that the best expert has an initial missing mass that is much larger than the number of experts – the waiting time of GT-UCB is comparable to $T^*(\alpha')$, up to factor that is only logarithmic in the waiting time of the oracle strategy. α' is smaller than α – hence $T^*(\alpha')$ is larger than $T^*(\alpha)$ – by an offset that is inversely proportional to the initial missing mass of the worst expert. This essentially says that, if we deal with large graphs, and if the experts trigger reasonably large spreads, our algorithm is competitive with the oracle. In practice, we will see in the experimental section that GT-UCB consistently leads to high-quality spreads, even on smaller networks.

6 EXPERIMENTS

We conducted experiments on graphs used in the IM literature and a dataset crawled from Twitter. All methods are implemented¹ in C++ and simulations are done on an Ubuntu 16.04 machine with an Intel Xeon 2.4GHz CPU 20 cores and 98GB of RAM.

6.1 Classic datasets

Similarly to [19], we have tested our algorithm on three publicly available datasets. NetHEPT and DBLP are collaboration networks, where undirected edges are drawn between authors which have collaborated on at least one paper. HepPh is a citation graph, where

¹The code is available at <https://github.com/smaniu/oim>.




Figure 2: Impact of K and ρ on influence spread.

a directed edge is established when an author cited at least one paper of another author. The datasets are summarized in Table 1. We emphasize that we kept the datasets relatively small to allow for comparison with computation-heavy baselines – even though GT-UCB easily scales to large data.

Table 1: Summary of the datasets.

Dataset	NetHEPT	HepPh	DBLP	Twitter
# of nodes	15.2K	34.5K	317K	12.2M
# of edges	62.7K	422K	2.1M	40.1M

Diffusion models. In the work closest to ours, Lei et al. [19] compared their solution on the Weighted Cascade instance of IC, where the influence probabilities on incoming edges sum to 1. More precisely, every edge (u, v) has weight $1/d_v$ where d_v is the in-degree of node v . In this experimental study, and to illustrate that our approach is model-free, we added two diffusion scenarios to our set of experiments. First, we included the tri-valency model (TV) which associates randomly a probability 0.1, 0.01 or 0.001 to every edge and follows the IC propagation model. We also added experiments under the Linear Threshold (LT) model, where the edges probabilities are set like in the WC case and where thresholds on nodes are sampled uniformly from $[0, 1]$. All the diffusion models simulate the persistence property: subsequent influence cascade may pass through previously activated nodes.

Baselines. We compare GT-UCB to several baselines. RANDOM chooses a random node from the graph at each round. MAXDEGREE selects the node with the i -th largest out-degree at step i . Finally, EG corresponds to the confidence-bound explore-exploit method with exponentiated gradient update from [19], the state-of-the-art method for the OIMP problem. We use this baseline on WC and TV weighted graphs and tune parameters in accordance to the results of their experiments: Maximum Likelihood Estimation is adopted for graph update and edge priors are set to Beta(1, 20). These baselines are compared to an ORACLE that knows beforehand the model together with diffusion probabilities. At each round, it runs an IM approximated algorithm – PMC for IC propagation, SSA for LT. Note that previously activated nodes are not counted when estimating the value of a node with PMC or SSA, thus, making ORACLE an adaptive strategy.

GT-UCB parameters. We first analyze the effects of the different possible settings for GT-UCB. We show in Fig. 2b and 2d the impact of the expert extraction method ρ on HepPh and DBLP

under WC model. We observe that the spread is slightly affected by the extraction method: different datasets lead to different optimal ρ . On HepPh network, DivRank clearly leads to larger influence spreads. On DBLP, however, the extraction method has less impact on resulting cascades. We emphasize that on some other graph and model combinations – due to limited space, we refer the reader to the Appendix for more details – we observed that other extraction routines can perform better than DivRank. In summary, we note that GT-UCB performs consistently as long as the method leads to experts that are well spread over the graph. In the following, for each graph, we used the ρ algorithm with the best spread.

In Fig. 6b and 6d, we measure the impact of the number of experts K on the influence spread. We observe that, on DBLP, a small number of experts is sufficient to yield high-quality results. If too many experts (relative to the budget) are selected (e.g. $K = 200$), the initialization step required by GT-UCB is too long relative to the full budget, and hence GT-UCB does not reach its optimal spread – some experts still have a large missing mass at the end. On the other hand, a larger amount of experts leads to greater influence spreads on HepPh: this network is relatively small ($n = 34.5K$), and thus half of the nodes are already activated after 400 trials. By having more experts, we are able to access parts of the network that would not be accessible otherwise.

GT-UCB vs baselines. We evaluate the execution time of the different algorithms in Fig. 4. As expected, GT-UCB greatly outperforms EG (and ORACLE). The two baselines require the execution of an approximated IM algorithm at each round. In line with [1], we observed that SSA has prohibitive computational cost when incoming edge weights do not sum to one, which is the case with both WC and TV. Thus, both ORACLE and EG run PMC on all our experiments with IC propagation. GT-UCB is several orders of magnitude faster: it concentrates most its running time on extracting experts, while statistic updates and UCB computations are negligible.

In Fig. 3, we show the growth of the spread for GT-UCB and baselines. We can see that GT-UCB results in good quality spreads across every combination of network and diffusion model. Interestingly, on smaller graphs NetHEPT and HepPh, we observe an increase in the slope of spread after initialization (e.g. for $t = 100$ on HepPh with WC). This corresponds to the step when GT-UCB starts to select experts maximizing $b_k(t)$ in the main loop. It shows that our strategy adapts well to the previous activations, and chooses good experts at each iteration. EG performs well on NetHEPT and HepPh, especially under TV weight assignment. However, it fails to provide competitive cumulative spreads on DBLP. We believe

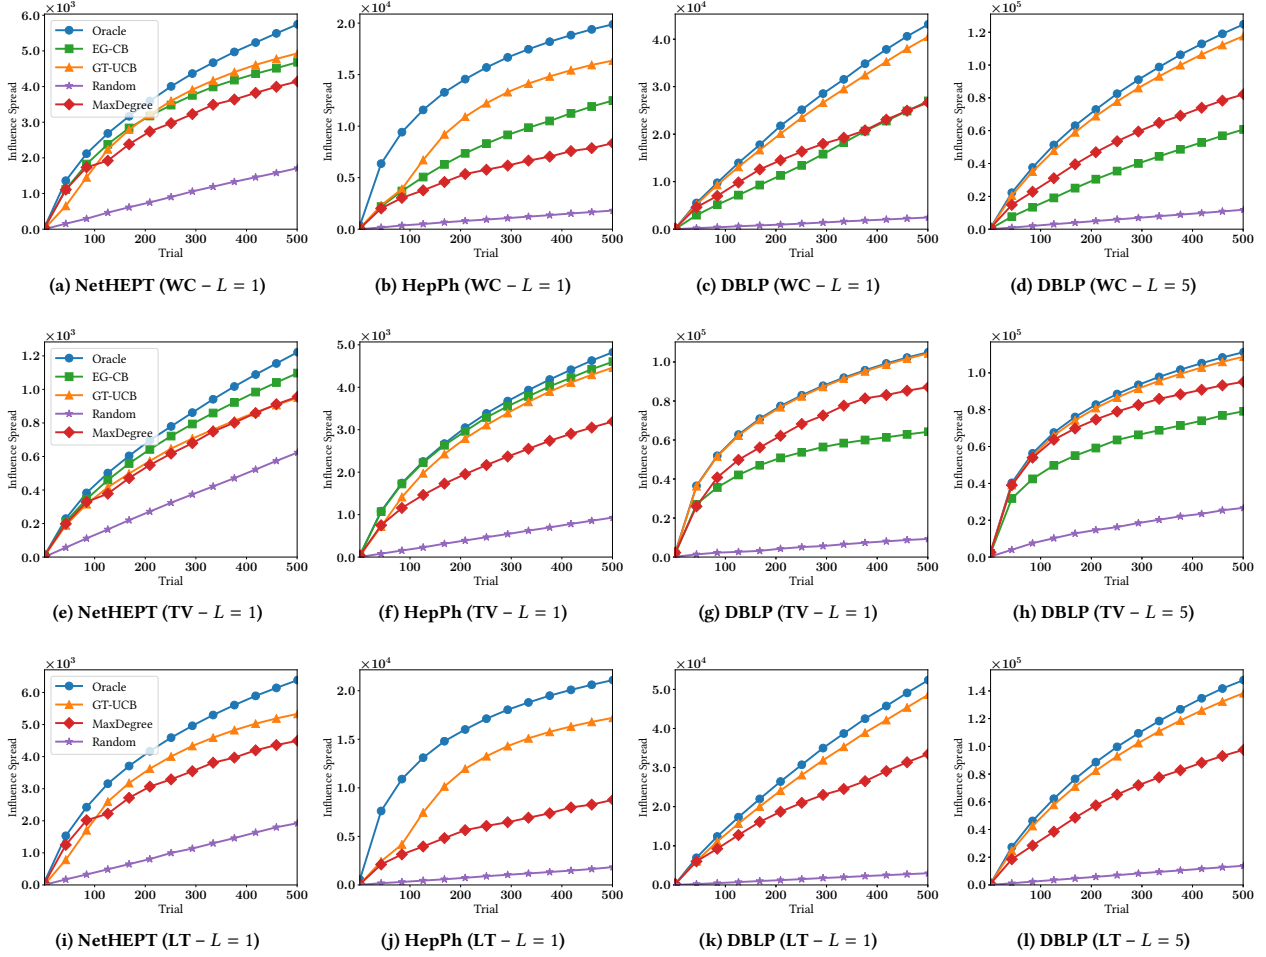


Figure 3: Growth of spreads against the number of rounds.

that EG tries to estimate too many parameters for a horizon $T = 500$. After reaching this time step, less than 10% of all nodes for WC, and 20% for TV, are activated. This implies that we have small confidence regarding many edge probability estimations as most nodes are located in parts of the graph that have never been explored.

6.2 Experiments on Twitter

We conclude the experimental section with simulations on Twitter data. We use a collection of tweets and re-tweets extracted via crawling in August 2012. For each original tweet, we find all corresponding retweets, and, for each user, we compute the empirical probability of a retweet occurring – this, in our case, is a proxy measure for influence probability. Specifically, for every user v “influenced” by u , i.e., v retweeted at least one original tweet from u – we compute the estimated diffusion probability: $p_{u,v} = \frac{|u\text{'s tweets retweeted by } v|}{|\text{tweets by } u|}$. On the left-hand side of Fig. 5, we show the histogram of resulting empirical probabilities. Unsurprisingly, we observe that most probabilities are very small – the 9th decile has value 0.045. This supports our initial motivation to rely on the missing mass to rapidly estimate the potential of a node in the graph. We also apply NPDC [26], a model-free approach to infer the underlying network using cascades of time intervals.




Figure 4: DBLP (WC) – Execution time.

We tested GT-UCB with different settings of K , and found that $K = 10$ experts provided the best results. This shows that the engagement of few influential people can suffice to conduct a marketing campaign. The chosen 10 experts have a very low intersection in their scopes of influence: the largest Dice coefficient of common neighbours is 0.07, supporting our theoretical choice of experts with non-overlapping support. In total, the 10 experts cover 810K users in our dataset. To test a realistic spread, a random cascade from the logs is chosen at every step. This provides realistic, model-free spread samples to the compared algorithms. Since Twitter only




Figure 5: (left) Empirical weights histogram (right) Twitter spread against rounds ($L = 1$).

contains successful re-tweets, we could not test against EG, which needs full activation feedback.

On the right-hand side of Fig. 5, we show the growth of the diffusion spread of GT-UCB against MAXDEGREE. We also compare ourselves to SINGLEHIGHESTDEGREE, a strategy playing the same most connected user at every step. We used this baseline, as we observed that MAXDEGREE was not competitive due to its rapidly diminishing influence properties. All baselines use the network inferred by NPDC. We notice that the single node strategy performs surprisingly well and MAXDEGREE extremely poorly; our GT-UCB, however, can increase the spread by around 10,000 users. This shows that relying on popular users for marketing campaigns can pay off; however, a good strategy should not target them indiscriminately, and should be able to adapt to the campaign feedback.

7 CONCLUSION

We have proposed an adaptive model-free approach to maximize the number of nodes activated in an arbitrary graph under the OIMP framework. From a social graph, the method requires an initial extraction of “expert nodes”. Subsequent online iterations are very fast, making it possible to scale to very large graphs where other approaches become infeasible. The efficiency of the GT-UCB algorithm comes from the fact that it only relies on an estimate of a single quantity for each expert node – its potential or missing mass. This novel approach was shown to be very competitive on classic benchmark tasks. As future work, we aim to work on the scenario where the persistence property is diminishing with time, i.e., where nodes may be less willing to relay information with subsequent activations.

REFERENCES

- [1] A. Arora, S. Galhotra, and S. Ranu. 2017. Debunking the Myths of Influence Maximization: An In-Depth Benchmarking Study. In *SIGMOD*. ACM.
- [2] N. Barbieri, F. Bonchi, and G. Manco. 2013. Topic-aware social influence propagation models. *Knowl. Inf. Syst.* 37, 3 (2013), 555–584.
- [3] D. Berend and A. Kontorovich. 2013. On the Concentration of the Missing Mass. In *Electronic Communications in Probability*. 1–7.
- [4] S. Boucheron, G. Lugosi, P. Massart, and M. Ledoux. 2013. *Concentration inequalities: a nonasymptotic theory of independence*. Oxford university press.
- [5] S. Bubeck and N. Cesa-Bianchi. 2012. Regret Analysis of Stochastic and Non-stochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122.
- [6] S. Bubeck, D. Ernst, and A. Garivier. 2013. Optimal discovery with probabilistic expert advice: finite time analysis and macroscopic optimality. *Journal of Machine Learning Research* 14, 1 (2013), 601–623.
- [7] W. Chen, T. Lin, Z. Tan, M. Zhao, and X. Zhou. 2016. Robust Influence Maximization. In *SIGKDD*. 795–804.
- [8] W. Chen, Y. Wang, Y. Yuan, and Q. Wang. 2016. Combinatorial Multi-armed Bandit and Its Extension to Probabilistically Triggered Arms. *JMLR* 17, 1 (Jan. 2016), 1746–1778.
- [9] N. Du, L. Song, H. Woo, and H. Zha. 2013. Uncover Topic-Sensitive Information Diffusion Networks. In *AISTATS*. 229–237.
- [10] D. Easley and J. Kleinberg. 2010. *Networks, Crowds, and Markets - Reasoning About a Highly Connected World*. Cambridge University Press.
- [11] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. 2011. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML*. 561–568.
- [12] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. 2012. Inferring Networks of Diffusion and Influence. *ACM Trans. Knowl. Discov. Data* 5, 4 (Feb. 2012), 21:1–21:37.
- [13] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf. 2013. Structure and dynamics of information pathways in online media. In *WSDM*. 23–32.
- [14] I. J. Good. 1953. The Population Frequencies of Species and the Estimation of Population Parameters. *Biometrika* 40, 3-4 (1953), 237.
- [15] A. Goyal, F. Bonchi, and L. Lakshmanan. 2010. Learning influence probabilities in social networks. In *WSDM*. 241–250.
- [16] P. Grabowicz, N. Ganguly, and K. Gummadi. 2016. Distinguishing between Topical and Non-Topical Information Diffusion Mechanisms in Social Media. In *ICWSM*. 151–160.
- [17] X. He and D. Kempe. 2016. Robust Influence Maximization. In *SIGKDD*. 885–894.
- [18] D. Kempe, J. Kleinberg, and É. Tardos. 2003. Maximizing the Spread of Influence Through a Social Network. In *SIGKDD*. ACM, 137–146.
- [19] S. Lei, S. Maniu, L. Mo, R. Cheng, and P. Senellart. 2015. Online Influence Maximization. In *SIGKDD*.
- [20] D. McAllester and L. Ortiz. 2003. Concentration Inequalities for the Missing Mass and for Histogram Rule Error. *JMLR* 4 (2003), 895–911.
- [21] D. McAllester and R. Schapire. 2000. On the Convergence Rate of Good-Turing Estimators. In *COLT*. 1–6.
- [22] Q. Mei, J. Guo, and D. Radev. 2010. DivRank: The Interplay of Prestige and Diversity in Information Networks. In *SIGKDD*.
- [23] H. T. Nguyen, M. T. Thai, and T. N. Dinh. 2016. Stop-and-Stare: Optimal Sampling Algorithms for Viral Marketing in Billion-scale Networks. In *SIGMOD*.
- [24] N. Ohsaka, T. Akiba, Y. Yoshida, and K. Kawarabayashi. 2014. Fast and Accurate Influence Maximization on Large Networks with Pruned Monte-Carlo Simulations. In *AAAI*.
- [25] D. Romero, B. Meeder, and J. Kleinberg. 2011. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *WWW*. 695–704.
- [26] Y. Rong, Q. Zhu, and H. Cheng. 2016. A Model-Free Approach to Infer the Diffusion Network from Event Cascade. In *CIKM*.
- [27] K. Saito, R. Nakano, and M. Kimura. 2008. Prediction of Information Diffusion Probabilities for Independent Cascade Model. In *KES*. 67–75.
- [28] Y. Tang, Y. Shi, and X. Xiao. 2015. Influence Maximization in Near-Linear Time: A Martingale Approach. In *SIGMOD*. 1539–1554.
- [29] Y. Tang, X. Xiao, and Y. Shi. 2014. Influence Maximization: Near-Optimal Time Complexity Meets Practical Efficiency. In *SIGMOD*. 75–86.
- [30] S. Vaswani, V.S. Lakshmanan, and M. Schmidt. 2015. Influence Maximization with Bandits. In *Workshop NIPS (NIPS '15)*.
- [31] S. Wang, X. Hu, P. Yu, and Z. Li. 2014. MMRate: inferring multi-aspect diffusion networks with multi-pattern cascades. In *SIGKDD*. 1246–1255.
- [32] D. Watts. 2003. *Six Degrees: The Science of a Connected Age*. W. W. Norton, NY.
- [33] D. Watts and P. Dodds. 2007. Influentials, networks, and public opinion formation. *Journal of Consumer Research* 34, 4 (2007), 441–458.
- [34] Z. Wen, B. Kveton, and M. Valko. 2016. Influence Maximization with Semi-Bandit Feedback. In *Working paper*.

A USEFUL LEMMAS

LEMMA A.1 (BENNETT'S INEQUALITY (THEOREM 2.9 AND 2.10 [4])). *Let X_1, \dots, X_n be independent random variables with finite variance such that $X_i \leq b$ for some $b > 0$ for all $i \leq n$. Let $S := \sum_{i=1}^n (X_i - \mathbb{E}[X_i])$ and $v := \sum_{i=1}^n \mathbb{E}[X_i^2]$. Writing $\phi(u) = e^u - u - 1$, then for all $t > 0$,*

$$\log \mathbb{E} \left[e^{tS} \right] \leq \frac{v}{b^2} \phi(bt) \leq \frac{vt^2}{2(1-bt/3)}.$$

This implies that, $\mathbb{P} \left(S > \sqrt{2v \log 1/\delta} + \frac{b}{3} \log 1/\delta \right) \leq \delta$.

LEMMA A.2 (LEMMA 7 - [3]). *Let $n \geq 1, \lambda \geq 0, p \in [0, 1]$ and $q = (1-p)^n$. Then,*

$$qe^{\lambda p(1-q)} + (1-q)e^{-\lambda pq} \leq \exp(p\lambda^2/(4n)) \quad (3)$$

$$qe^{\lambda p(q-1)} + (1-q)e^{\lambda pq} \leq \exp(p\lambda^2/(4n)) \quad (4)$$

B ANALYSIS OF THE WAITING TIME OF GOOD-UCB ALGORITHM

LEMMA B.1. *For any $s \geq 3$, $\mathbb{P} \left(\hat{R}_s \leq \hat{R}_{s-1} - \frac{\lambda}{e(s-2)} - \sqrt{\frac{2\lambda}{s-1} \log(1/\delta)} - \frac{1}{3(s-1)} \log(1/\delta) \right) \leq \delta$.*

PROOF. Denote by $X_s(x) := \frac{U_{s-1}(x)}{s-1} - \frac{U_s(x)}{s} \leq \frac{1}{s-1}$. We can rewrite $\hat{R}_{s-1} - \hat{R}_s = \sum_{x \in A} X_s(x)$ and can easily verify that

$$v(x) := \mathbb{E} [X_s(x)^2] = p(x)(1-p(x))^{s-2} \left(\frac{1}{s-1} - \frac{1-p(x)}{s} \right) \leq \frac{p(x)}{s-1}. \quad (5)$$

Let $t > 0$. By applying Lemma A.1, one obtains

$$\mathbb{P} \left(\hat{R}_{s-1} - \hat{R}_s \geq \mathbb{E} [\hat{R}_{s-1} - \hat{R}_s] + \sqrt{\frac{2\lambda}{s-1} \log(1/\delta)} + \frac{1}{3(s-1)} \log(1/\delta) \right) \leq \delta.$$

We conclude remarking that $\mathbb{E}[X_s(x)] = p(x)^2(1-p(x))^{s-2} \leq \frac{p(x)}{e(s-2)}$, that is, $\mathbb{E}[\hat{R}_{s-1} - \hat{R}_s] \leq \frac{\lambda}{e(s-2)}$. \square

THEOREM B.2 (STOPPING TIME). *Denote $\lambda^{\min} := \min_{k \in [K]} \lambda_k$ and $\lambda^{\max} := \max_{k \in [K]} \lambda_k$. Assume that $\lambda^{\min} \geq 13$. Then, for any $\alpha \in [\frac{13}{\lambda^{\min}}, 1]$, if we define $\tau^* := T^* \left(\alpha - \frac{13}{\lambda^{\min}} \right)$, with probability at least $1 - \frac{2K}{\lambda^{\max}}$,*

$$T_{UCB}(\alpha) \leq \tau^* + K\lambda^{\max} \log(4\tau^* + 11K\lambda^{\max}) + 2K.$$

PROOF. Let us define the following confidence bounds:

$$\begin{aligned} b_{k,s}^+(t) &:= (1 + \sqrt{2}) \sqrt{\frac{3\lambda_k \log(2t)}{s}} + \frac{\log(2t)}{s}, \\ b_{k,s}^-(t) &:= (1 + \sqrt{2}) \sqrt{\frac{3\lambda_k \log(2t)}{s}} + \frac{\log(2t)}{s} + \frac{\lambda_k}{s}, \text{ and} \\ c_{k,t}^-(t) &:= \frac{\lambda}{e(s-2)} + \sqrt{\frac{6\lambda_k \log(t)}{s-1}} + \frac{\log(t)}{s-1}. \end{aligned}$$

Let $S > 0$. Using these definitions, we introduce the following events:

$$\begin{aligned} \mathcal{F} &:= \left\{ \forall k \in [K], \forall t > S, \forall s \leq t, \hat{R}_{k,s} - b_{k,s}^-(t) \leq R_{k,s} \leq \hat{R}_{k,s} + b_{k,s}^+(t) \right\}, \\ \mathcal{G} &:= \left\{ \forall k \in [K], \forall s \geq S, \hat{R}_{k,s} \geq \hat{R}_{k,s-1} - c_{k,s}^-(t) \right\}, \\ \mathcal{E} &:= \mathcal{F} \cap \mathcal{G}. \end{aligned}$$

Using Theorem 3.5, Lemma B.1 and a union bound, one obtains $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{2K}{S}$ (by setting $\delta \equiv \frac{1}{t^3}$). Indeed,

$$\mathbb{P}(\bar{\mathcal{E}}) \leq \mathbb{P}(\bar{\mathcal{F}}) + \mathbb{P}(\bar{\mathcal{G}}) \leq 2 \sum_{k=1}^K \sum_{t>S} \sum_{s \leq t} \frac{1}{t^3} = 2K \sum_{t>S} \frac{1}{t^2} \leq \frac{2K}{S}.$$

In the following, we work on the event \mathcal{E} . Recall that we want to control $T_{UCB}(\alpha)$, the time at which every expert attains a missing mass smaller than α following Good-UCB strategy. We aim at comparing $T_{UCB}(\alpha)$ to $T^*(\alpha)$, the same quantity following the omniscient strategy. With that in mind, one can write:

$$T_{UCB}(\alpha) = \min \{t : \forall k \in [K], R_{k, N_k(t)} \leq \alpha \lambda_k\},$$

$$T^*(\alpha) = \sum_{k=1}^K T_k^*(\alpha), \text{ where } T_k^*(\alpha) = \min \{s : R_{k,s} \leq \alpha \lambda_k\}.$$

Following ideas from [6], we will control $T_{UCB}(\alpha)$ by comparing it to $U(\alpha)$ defined below, and which replaces the missing mass by an upper bound on the *estimator* of the missing mass (the Good-Turing estimator). Indeed, remind that we can control this on event \mathcal{F} .

$$U(\alpha) = \min \left\{ t \geq 1 : \forall k \in [K], \hat{R}_{k, N_k(t)} + b_{k, N_k(t)}^+(t) \leq \alpha \lambda_k \right\}.$$

Let $S' \geq S$. On event \mathcal{E} , one has that $T_{UCB}(\alpha) \leq \max(S', U(\alpha))$. If $U(\alpha) \geq S'$, one has

$$\begin{aligned} R_{k, N_k(U(\alpha))} &\geq \hat{R}_{k, N_k(U(\alpha))} - b_{k, N_k(U(\alpha))}^-(U(\alpha)) && \text{(we are on event } \mathcal{F} \text{ and } U(\alpha) > S' \geq S) \\ &\geq \hat{R}_{k, N_k(U(\alpha))-1} - b_{k, N_k(U(\alpha))}^-(U(\alpha)) - c_{k, N_k(U(\alpha))}^-(U(\alpha)) && \text{(where are on event } \mathcal{G}) \\ &\geq \left(\alpha \lambda_k - b_{k, N_k(U(\alpha))-1}^+(U(\alpha)) \right) - b_{k, N_k(U(\alpha))}^-(U(\alpha)) - c_{k, N_k(U(\alpha))}^-(U(\alpha)) \end{aligned}$$

The third inequality's justification is more evolved. Let t be the time such that $N_k(t) = N_k(U(\alpha)) - 1$ and $N_k(t+1) = N_k(U(\alpha))$. This implies that k is the chosen expert at time t , that is, the one maximizing the Good-UCB index. Moreover, since $t < U(\alpha)$, one knows that this index is greater than $\alpha \lambda_k$.

If $N_k(U(\alpha)) \geq S' + 2$, some basic maths calculations lead to

$$R_{k, N_k(U(\alpha))} \geq \alpha \lambda_k - 11 \sqrt{\frac{\lambda_k \log(2U(\alpha))}{S'}} - \frac{3 \log(2U(\alpha))}{S'} - \frac{3 \lambda_k}{2S'}$$

We denote by $\lambda^{max} := \max_k \lambda_k$. If we take $S' = \lambda^{max} \log(2U(\alpha))$, we can rewrite previous inequality as

$$R_{k, N_k(U(\alpha))} \geq \alpha \lambda_k - 11 - \frac{3}{\lambda^{max}} - \frac{3}{2}$$

Thus, by definition of $T_k^*(\alpha)$, and if $\lambda^{max} > 6$, one gets

$$N_{k, U(\alpha)} \leq T_k^* \left(\alpha - \frac{13}{\lambda_k} \right) + S' + 2.$$

Finally, if we denote by $\lambda^{min} = \min_k \lambda_k$, we obtain that

$$U(\alpha) \leq K(S' + 2) + T^* \left(\alpha - \frac{13}{\lambda^{min}} \right).$$

We now apply Lemma B.3. We obtain that

$$U(\alpha) \leq 2K + \tau^* + K \lambda^{max} \log(8K + 4\tau^* + 10K \lambda^{max}) \leq \tau^* + K \lambda^{max} \log(4\tau^* + 11K \lambda^{max}) + 2K.$$

We conclude with $T_{UCB}(\alpha) \leq \max(S', U(\alpha))$. □

LEMMA B.3 (LEMMA 3 FROM [6]). *Let $a > 0$, $b \geq 0.4$, and $x \geq e$, such that $x \leq a + b \log x$. Then one has*

$$x \leq a + b \log(2a + 4b \log(4b)).$$

Moreover, we add that if $b \geq 3$, then $x \leq a + b \log(2a + 5b)$.

C ADDITIONAL FIGURES

We present in this section several figures we could not include in the main paper to to the lack of space. We show in Fig. 6 how the number of experts K impacts the quality of GT-UCB spreads accross all network and diffusion setting combinations. In Fig. 7, we display how the expert extraction routine ρ impacts the quality of spreads of our algorithm.

We show in Fig. 8 the Dice and Jaccard similarity matrices of the 10 experts extracted by GT-UCB in our experiments.

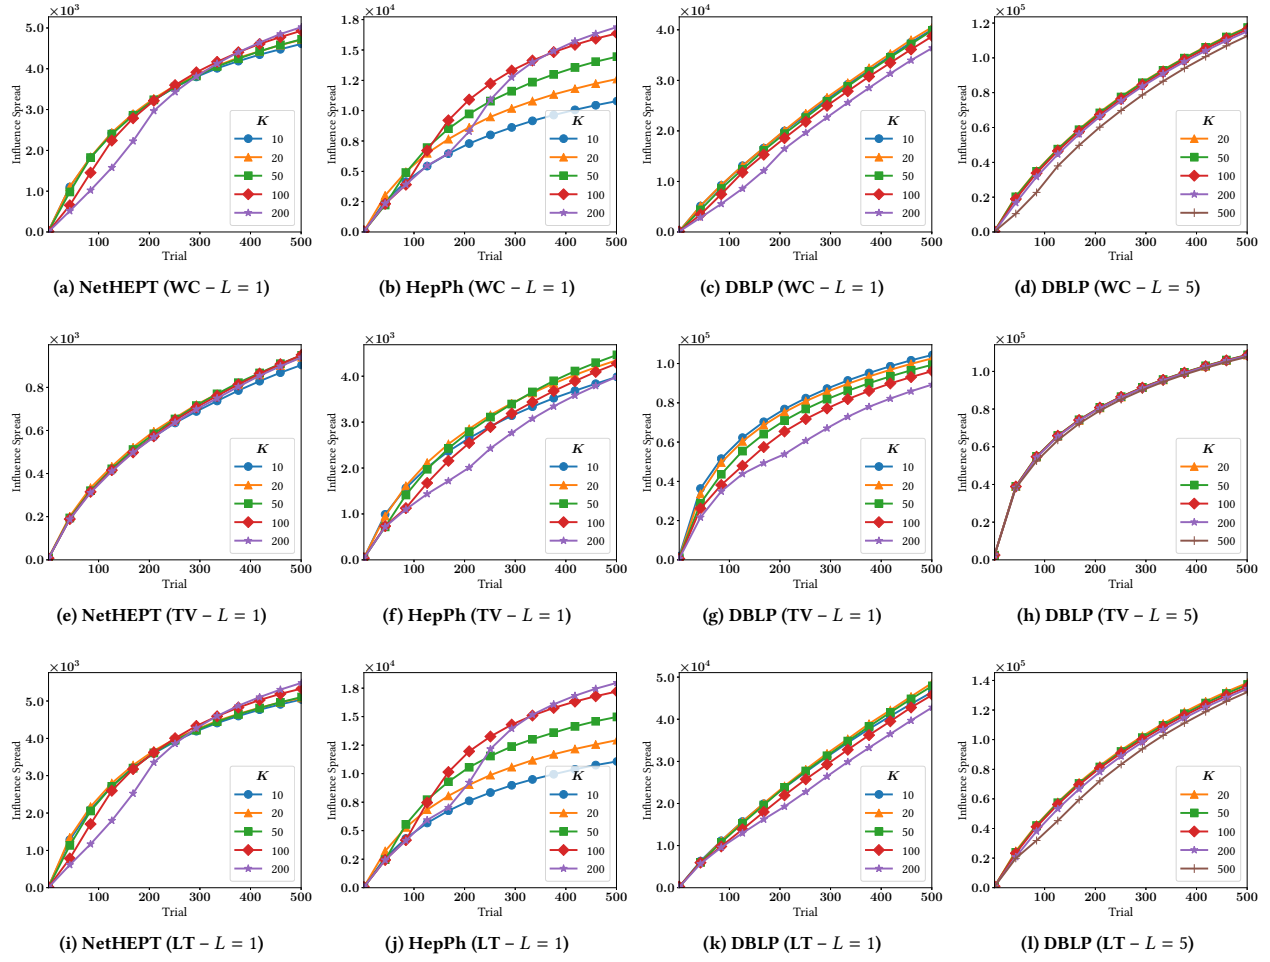


Figure 6: Impact of K on influence spread.

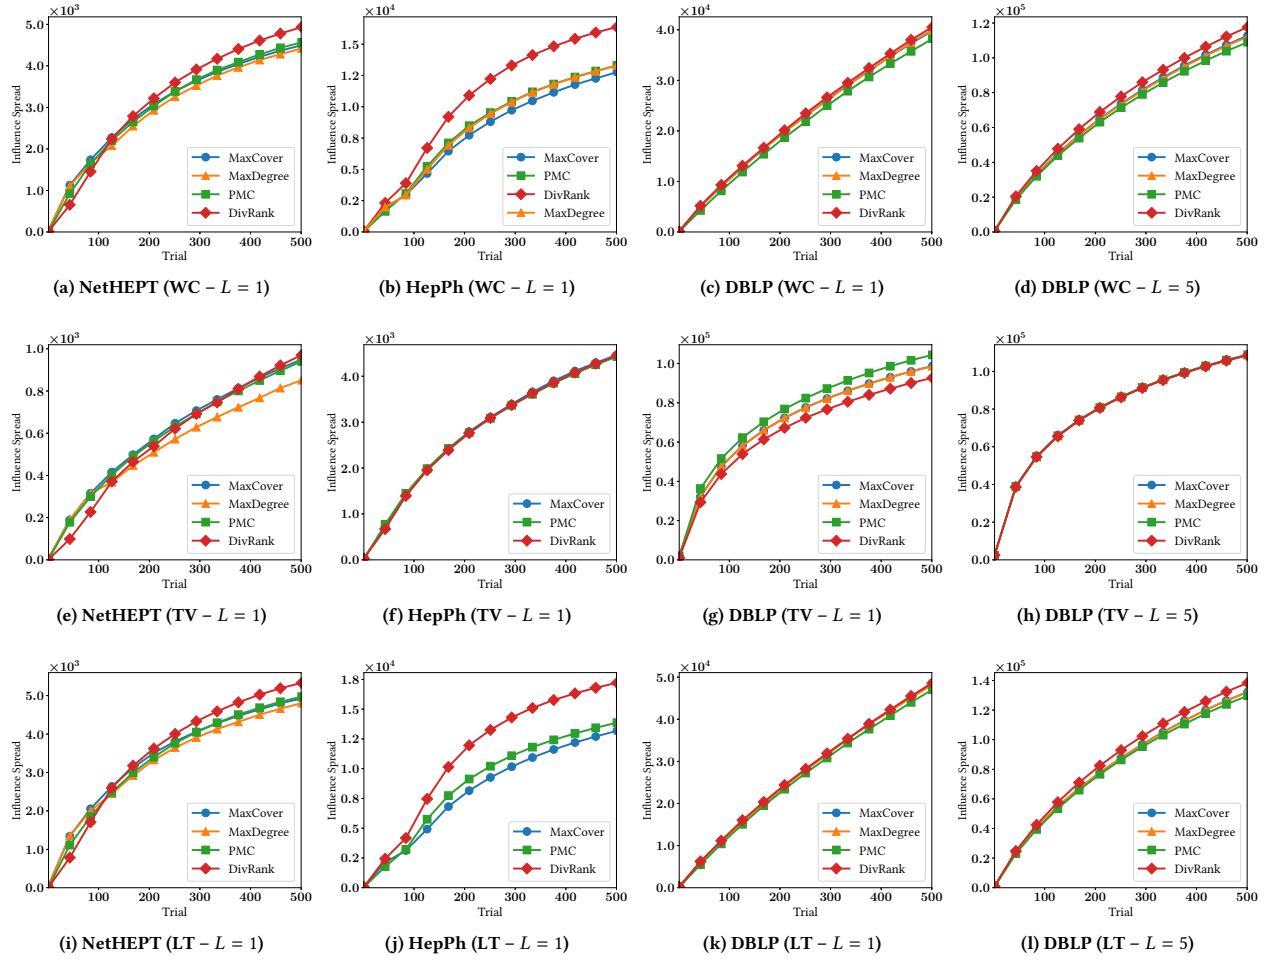


Figure 7: Impact of K on influence spread.




Figure 8: Dice (left) and Jaccard (right) matrices for 10 experts